

**U.S. DEPARTMENT OF COMMERCE
National Technical Information Service**

AD-A032 229

The Use of Spectral Analytic Techniques in Economics

Rund Corp Santa Monica Calif

Aug 72

328
ADA032229

REPRODUCED BY
NATIONAL TECHNICAL
INFORMATION SERVICE
U. S. DEPARTMENT OF COMMERCE
SPRINGFIELD, VA. 22161

THE USE OF SPECTRAL ANALYTIC TECHNIQUES IN ECONOMICS

Richard V.L. Cooper*

The Rand Corporation, Santa Monica, California

ACCESSION No.	
NTIS	White Section <input checked="" type="checkbox"/>
DOC	Buff Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION	
BY	
DISTRIBUTION/AVAILABILITY CODES	
DOC.	AVAIL. and/or SPECIAL
A	

I. INTRODUCTION

In recent years spectral analysis has become an increasingly popular tool for use in the estimation of econometric models. The interest in spectral analysis stems largely from the desire to have additional methodologies for estimating economic time series relationships when the more traditional time-domain methods are not satisfactory. However, there has been relatively little discussion on the advisability of using spectral analysis for such purposes.

In this paper I comment upon the applicability of spectral analysis to econometric problems and provide some guidelines as to when spectral analysis may prove a useful tool for economists. Since knowledge of spectral analysis generally requires a substantial investment of time, I first outline its possible uses so that the economist unfamiliar with the techniques can better evaluate his potential need. An outline of the types of economic problems particularly susceptible to analysis by spectral techniques should interest the investigator currently dealing with those types of problems.

* Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The Rand Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The Rand Corporation as a courtesy to members of its staff.

This paper is primarily concerned with the frequency-domain applications of spectral analysis and the advisability of using such analysis for examining economic time series relationships. That is, I do not dwell upon the techniques developed by Lannan (1965) for using spectral estimates to estimate distributed lag models. These techniques are really just an alternative method for estimating time-domain parameters. One of the major stumbling blocks in the application of spectral analysis in economics is the interpretation of the frequency domain.

In Section II I provide a brief description of spectral analytic techniques. In Section III theoretical and practical reasons for and against the use of spectral analysis in economics are discussed. As an illustration of the rationale for using spectral analysis in economics, I discuss three specific examples from the literature in Section IV. Finally, Section V presents the conclusions.

Spectral analysis is found to be a useful tool for a limited set of economic hypotheses. The major conceptual problem in the application of spectral analysis to economics is the economic interpretation of frequency. Therefore, a major requirement to be satisfied is that there be a reasonable frequency-domain interpretation of the hypothesis. On the other hand, a number of practical problems may seriously limit the use of spectral analysis in econometrics. The most important of these are the large data requirements necessary for estimation, the limited number of variables that can reasonably be included in the model, and the matter of how to interpret the large number of estimates for any given model (or alternatively, the lack of summary statistics). Consequently, although spectral analysis may prove to be a useful tool, its applicability (in the frequency domain) is limited.

II. DESCRIPTION OF SPECTRAL ANALYSIS

Since there are a number of excellent references dealing with the presentation of spectral analysis,¹ my description of spectral techniques will be brief. However, it would be useful at least to summarize the methodology to put the discussion into the proper perspective.

Spectral analysis is a method by which time series data are converted into the frequency domain for examination. Through the Fourier transform it can be shown that any stationary time series can be decomposed into a summation of sinusoidal waves of different frequencies (or equivalently, different periods). Each of these waves is completely described by its frequency, amplitude, and phase shift. The frequency is simply the fraction of a cycle that is completed in one period. The amplitude is the height of the wave. The phase shift is the fraction of a cycle that the wave is displaced from zero. Therefore, given a real valued stationary time series x_t ($t = 1, 2, \dots, T$), where x_t are deviations from the mean, x_t may be written:

$$x_t = 2 \sum_{k=1}^{n-1} A_k \cos(2\pi k/t + \varphi_k) + A_n \cos(\pi),$$

where

$n = T/2$ (suppose T even),

A_k = amplitude for frequency k , and

φ_k = phase shift for frequency k .

¹For example, see Jenkins and Watts (1968), Fishman (1970), among others.

Figure 1 demonstrates the decomposition of a continuous signal into the summation of three different cosine curves. Figure 1a shows the continuous signal; Figure 1b shows the decomposition. It is clear from Figure 1b that most of the variance in the observed signal is described by the cosine with the lowest frequency (that is, the longest period).

Expression (1) may be rewritten:

$$(2) \quad x_t = \sum_{m=-n}^{n-1} X_m e^{i2\pi mt/T}$$

where

$$e^{i\theta} = \cos \theta + i \sin \theta, \quad i = \sqrt{-1}.$$

On the other hand, the theory of Fourier transforms states that X_m may be written as an inverse Fourier transform of x_t :

$$(3) \quad X_m = \frac{1}{T} \sum_{t=-n}^{n-1} x_t e^{-i2\pi mt/T}.$$

Expressions (2) and (3) indicate that a time series may be viewed either through the time domain $\{x_t\}$ or through the frequency domain $\{X_f\}$. That is, $\{x_t\}$ and $\{X_f\}$ are equivalent; they just represent alternative views of the same data. Although economists are generally quite familiar with the time domain, many are not familiar with the frequency domain. Therefore the meaning of the frequency domain is probably worth reiteration. The time domain refers to the description of data in terms of their values over time. The frequency domain refers to the same data, but with reference to the collection of sine and cosine waves that would be required to produce those data.

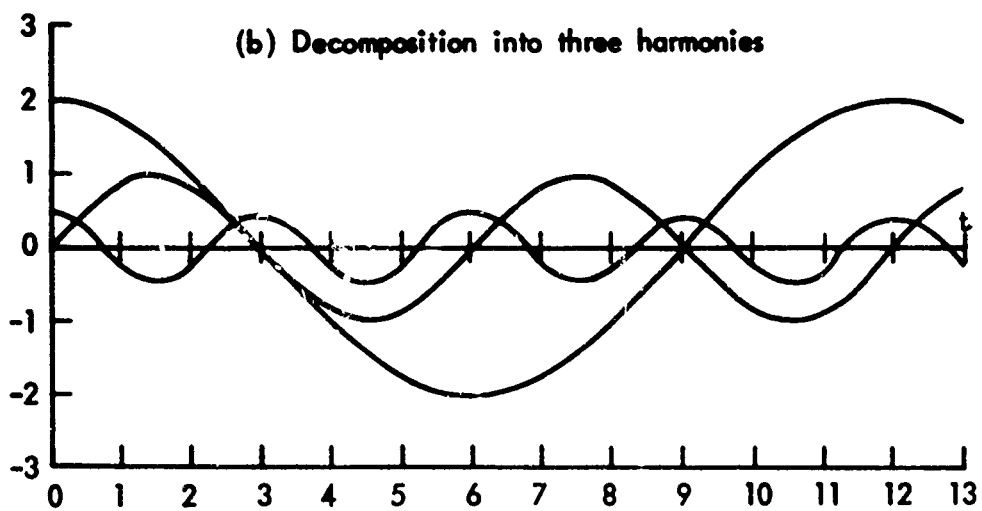
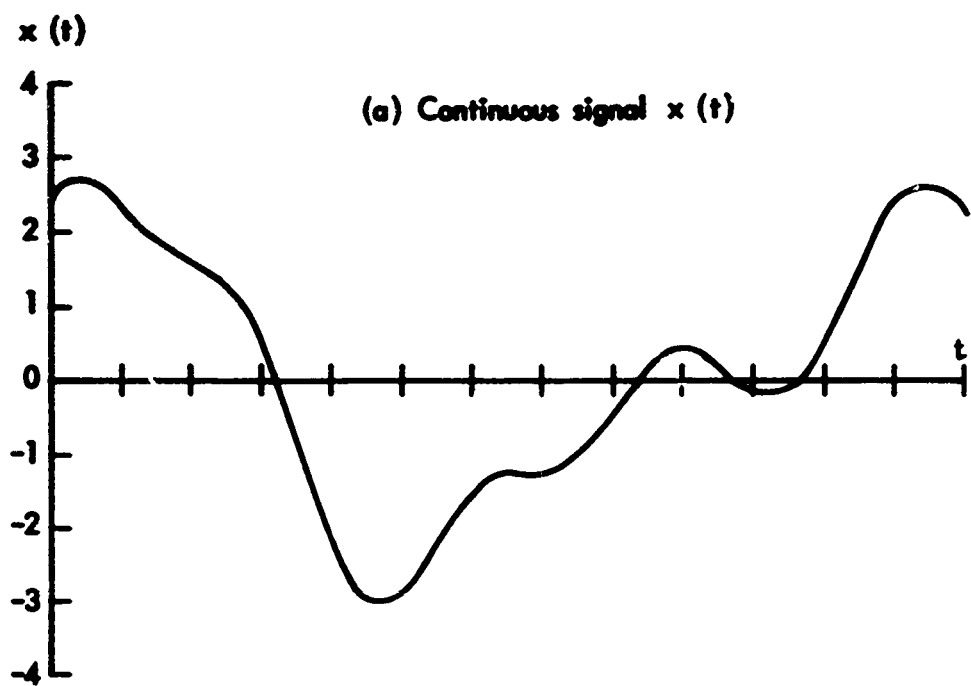


Fig. 1 — Signal and its decomposition

The functional relationship between X_f and f is known as the spectrum (or power spectrum). The spectrum of a time series represents a decomposition of variance in the frequency domain (versus the time-domain decomposition of variance techniques usually used by economists). In other words, the spectrum measures the relative contribution of different frequency bands to the variance of the time series. Those frequencies that contribute most to the variation in the time series will have the largest values for the power spectrum; those frequencies that contribute least to the variation will have the smallest values. For example, the spectrum of a time series dominated by long-run swings (for example, trends and long cycles) will have most of its power in the low frequency bands, since long cycles primarily contribute to the variance in that series. Alternatively, a time series dominated by short-run fluctuations will have a spectrum with most of the power in the high frequency bands.

Since the "true" spectrum is not observed, it must be estimated. It can be shown that the spectrum of the autocovariance function for a time series x_t is equivalent to the spectrum of x_t itself. However, using the entire estimated autocovariance function (all T terms) does not produce a consistent estimator of the spectrum, for the analyst never leaves the small sample situation. On the other hand, a consistent estimator of the spectrum does result from Fourier transforming a truncated autocovariance function for the first M terms, where $M < T$. The estimator can be further improved by averaging the spectrum estimate for a given frequency with the spectrum estimates for some nearby frequencies. The weighting scheme for this averaging is known as the

spectral window, or averaging kernel. Therefore, by Fourier transforming a truncated autocovariance function through a spectral window the analyst is able to obtain consistent estimates of the spectrum.

The discussion up to this point has dealt with univariate spectral analysis. Just as two series may be studied jointly in the time domain, they may be studied jointly in the frequency domain by means of cross-spectral analysis. In cross-spectral analysis the frequencies of one variable are compared with their counterparts of the other variable. Therefore, for each estimated frequency, the analyst can examine the relationship between the two variables.

The input to the cross spectrum is the cross-covariance function (just as the input to the power spectrum is the autocovariance function). The cross spectrum is the Fourier transform of the cross-covariance function. Since the cross covariance is not, in general, an even function, the cross spectrum will generally be complex valued.¹ In addition to being complex valued, the cross spectrum has no readily apparent interpretation (just as the cross-covariance function has no ready interpretation). However, three very useful measures do arise from the cross spectrum and the two sets of power spectra. First, the coherence provides a measure of the correlation between the first variable and the second for each given frequency. The gain measures the relation between the amplitude of the explanatory variable and the amplitude of the cross spectrum for each given frequency. Third, the phase spectrum measures the delay between the two variables for each given frequency.

¹The power spectrum is real valued because the autocovariance function is even ($R(u) = R(-u)$).

Many of the functions in the frequency domain have direct analogs in the ordinary regression (time domain) model. The coherence is analogous to the R^2 measure in the regression model. The gain is much like the regression coefficient. The phase shift is similar to a lag operator in the regression model, showing the delay between the two time series.¹

Finally, multivariate spectral analysis can be used to examine the relationship between a dependent variable and several explanatory variables, just as multiple regression analysis allows examination of the relationship between one dependent variable and several explanatory variables in the time domain. However, the usefulness of multivariate spectral analysis in economics is likely to be very limited because of the extremely large number of parameters necessary for estimation of the particular spectra.

In summary, spectral analysis provides a description of time series variables in the frequency domain. That is, since any stationary time series can be decomposed into a collection of sine and cosine waves of different frequencies (via the Fourier transform), that time series can be examined across frequencies, not just across time.

¹Some care is needed in interpreting the phase shift in the time domain, as shown by Hause (1971). The phase shift is analogous to a lag operator. Yet one cannot directly translate pure delay in the frequency domain into distributed lag models. See Hause for further explanation.

III. SPECTRAL ANALYSIS IN ECONOMICS

With the numerous methods for investigating the relationships between economic variables in the time domain, why should one wish to use spectral analysis? It is the thesis of this paper that spectral analysis is useful in some instances when time-domain techniques do not prove adequate. Certain types of hypotheses may be better examined in the frequency domain than in the time domain. The special characteristics that describe these hypotheses are discussed below. The practical difficulties that limit the usefulness of spectral analysis are also discussed.

GENERAL COMMENTS

It has been suggested that one area where spectral analysis might be useful in economics is in the estimation of distributed lag models when the analyst has little or no prior information as to the shape of the distributed lag or even the number of terms to be included in the estimation. This problem can be troublesome in the time domain because of the multicollinearity between the lagged variables. When additional lags are introduced into a least squares estimation, the estimates for the previously estimated coefficients change -- often markedly. That is, the estimate for any one coefficient depends partly upon the number of terms included in the regression.

Spectral analysis has been proposed as a method for avoiding this problem, since it does not impose prior restrictions on the model -- the data is said "to speak for itself." However, although the phase spectrum may be useful for indicating the direction of leads and lags, Hause (1971) notes that it is not appropriate to view the phase spectrum

in the framework of a distributed lag model. The phase represents pure delay in the frequency domain, not a distributed lag. Therefore, spectral analysis may not be useful in itself for the estimation of distributed lag models. On the other hand, it may be useful for identifying general leads and lags.

An alternative method for estimating distributed lags using spectral analysis is suggested by Hannan (1965). The spectral estimates are converted back into the time domain so that successive coefficients are estimated independently of one another. This technique seems promising, but this independence of the estimates comes at the cost of possibly greater variance. Even so, Hannan's technique may be the best way out of a troublesome problem. Hannan's technique is not dealt with in this paper because, as noted in the introduction, this paper is concerned with frequency-domain applications of spectral analysis. Hannan's technique represents an alternative way of estimating the time-domain parameters and, therefore, should be considered in that framework.

Perhaps the greatest conceptual difficulty faced by economists in the application of spectral analysis is working in the frequency domain rather than in the time domain. Most economic hypotheses are given in terms of time, not frequency. It could be argued that virtually no economic hypotheses are given in terms of frequency. However, some economic hypotheses are given in terms of separate components, where each component is expected to vary predictably over some regular period of time; hence, these types of hypotheses may have some meaning in the frequency domain. For example, consider the model in which the time

series is postulated to consist of trend, cyclical, seasonal, and random components. The trend component is expected to vary only over long periods of time; the cyclical component is expected to vary somewhat predictably over the course of a business cycle; the seasonal component is expected to vary predictably over the course of a year; and the random component has no predictable behavior. Although models of this sort are given in terms of time periods, they may also be thought of in frequency terms.

Again, because economic hypotheses are seldom thought of in frequency terms the application of spectral analysis in economics has been limited. It could probably be argued that not only are most economic hypotheses not given in frequency terms but also that it is not even appropriate to think of most economic hypotheses in terms of frequency. Indeed, it could be further argued that it is not even appropriate to think of component models like the one described above in the frequency domain, for frequency-domain analysis imposes very restrictive conditions on the meaning of frequency. That is, spectral analysis views frequencies in terms of a narrow frequency band that describes a set of cosine waves. Economists seldom attach such a precise meaning to the components of a model, with the possible exception of the seasonal component, which is expected to adhere to a strict yearly cycle.¹

The above comments seem to imply that spectral analysis is never applicable to economics, but that was not intended, for it will be argued

¹It is difficult to argue for a precise periodic interpretation even for the seasonal component since there are leap years, different numbers of days per month, and holidays. It is partly for this reason that spectral analysis is far more useful than periodogram analysis, for the spectrum is actually an average of nearby frequency bands.

below that spectral analysis may actually be the most appropriate tool to use in the estimation and testing of some types of economic hypotheses. Instead, these comments actually apply strictly only to harmonic (or periodogram) analysis, as used about the turn of the century.¹ Since spectral analysis provides a measure of the average contribution to variation of a frequency band (rather than one single frequency), spectral analysis has more meaning in economic than strict harmonic analysis. However, the investigator must still keep in mind that the frequency domain will often have no real meaning for an economic hypothesis. Therefore, the distinction between the time domain and the frequency domain must be clear. This does not preclude thinking of frequencies as their period counterparts. In fact, in most cases it is probably useful for economists to think of cycle lengths, rather than the frequencies corresponding to those cycles. Again, this is largely because most economic hypotheses are stated in terms of time.

CRITERIA FOR THE USE OF SPECTRAL ANALYSIS IN ECONOMICS

The initial criterion that an economic hypothesis must satisfy for spectral analysis to be of much use is that it have some plausible interpretation in the frequency domain. Since the spectral results must ultimately be interpreted in the framework of economics, it is essential that the hypothesis have some meaning in the frequency domain.

Hypotheses that are amenable to spectral analysis fall into two broad classes: (1) those that are perhaps best stated in the frequency domain, and (2) those that are stated in the frequency domain to take

¹See Fishman (1970).

advantage of special properties of the spectral estimators. The first class of hypotheses consists of those types of models in which it is expected that the variables in the model have a reasonably well-defined periodic nature and that the variables are related to one another in a periodic sense. For example, suppose one wanted to test whether seasonal fluctuations in the demand for money produce seasonal fluctuations in interest rates. This class of hypotheses defines relationships between variables in a periodic sense and attempts to identify the parameters of the relationships.

It would be expected that the first class of hypotheses is fairly small. On the other hand, there may be a significant number of hypotheses that, although best stated in the time domain, are more susceptible to estimation or testing in the frequency domain. To understand why it might be preferable to use the frequency domain when the hypothesis itself is best stated in the time domain, one must keep in mind that the time domain (for example, the regression model) averages the relationships between variables over all frequencies. That is, the regression coefficient represents the effects among variables as an average of the effects over all frequencies. This can be seen by reviewing expression (3) in the previous section.

In most cases, the fact that time-domain descriptions of the data are averages across frequencies presents no real problems. However, if the investigator is interested in a hypothesis positing different relationships for different parts of the spectrum, then spectral analysis -- in the frequency domain -- may prove to be a more useful tool. Since time-domain analysis averages the relationship across all

frequencies, the investigator may not be able to separate differing behavior across different frequencies by time-domain techniques. On the other hand, spectral analysis permits explicit examination across frequencies.

Even if the analyst expects different behavior across different frequencies, averaging across frequency bands may not cause serious problems since most economic variables are dominated by a particular subset of frequencies. In fact, a majority of economic time series variables probably have most of their variation explained by the low frequency end of the spectrum -- that is, they are dominated by relatively long-run changes. Therefore, the behavior in the other parts of the spectrum has little effect on the total variation in the time series.

The above statement should be interpreted with some care. If, in fact, the hypothesis indicates differing behavior across different frequencies and if the time-domain model does not account for the differing behavior, then the time-domain model will misspecify and the coefficients will not be consistently estimated.¹ The effects of such a time-series estimation are perhaps most easily viewed through the errors-in-the-variables (EV) model. The analog to the "true" variable in the EV model for the frequency domain may be taken to be the set of frequencies that corresponds to the basic hypothesis (the "primary" frequencies). The fact that other frequencies are not expected to behave in the same manner corresponds to the "errors" in the EV model. Therefore,

¹An example of this problem is considered in the next section.

the smaller the amount of variance explained by these other frequencies, the smaller are the errors in the EV model; the smaller the amount of variance explained by these other frequencies, the less serious it will be to estimate the model in the time-domain even if the time-domain model does not account for differing behavior across different frequencies. Of course, if the time-domain model does account for the differing behavior across frequencies, then consistent time-domain estimators will result.¹

These comments imply that the usefulness of spectral analysis in economics is largely dependent upon (1) the plausibility of a frequency-domain interpretation for the model, (2) the consistency of behavior across frequencies, and (3) the amount of variance explained by the primary frequencies relative to the other frequencies.

PRACTICAL LIMITATIONS TO THE USE OF SPECTRAL ANALYSIS

In addition to the theoretical rationale for and against the use of spectral analysis in economics, one must also consider the practical limitations. Three are considered here: (1) the large amount of data required, (2) the interpretation of the phase angle, and (3) the lack of summary statistics.

Generally speaking, spectral analysis requires large amounts of data. The reason for the large amounts of data is directly related to the resolution often required in economics. The resolution of the spectral estimates is a function of the truncation point M : the larger M is, the finer the resolution. On the other hand, as M increases

¹That is, if the time-domain model is specified correctly, then one need not go to the frequency domain for estimation. An example of this type of problem is the permanent income hypothesis examined in the next section.

relative to the total sample size, the variance of the estimator increases. Therefore, some compromise must be made between resolution and variance. For a spectrum that is expected to contain narrow peaks or troughs, a small frequency bandwidth (large M) is needed to resolve the spectrum adequately. For a spectrum that is relatively flat, a much broader bandwidth may be chosen (smaller M).

This issue may be approached from another point of view. For the most commonly used spectral windows, the bias of the spectral estimator is a function of the second derivative of the spectrum¹ such that the bias is large when the absolute value of the second derivative is large (the size of the second derivative determines the sign of the bias). For a spectrum with narrow peaks or troughs -- hence, a nonnegligible second derivative -- the bias will be greater than for a relatively flat spectrum. Therefore, in order to avoid serious bias, a narrower bandwidth must be chosen.

The problem of resolution enters in a special way for economic studies. As noted earlier, it is often useful to think of the spectrum in economics not in terms of frequencies, but in terms of period lengths corresponding to frequencies. It is often the longer periods (lower frequencies) that are of special interest to economists. Although the spectrum is estimated at $M + 1$ frequencies (zero frequency plus the next M frequencies) equidistant from one another, the corresponding periods are not equidistant from one another. As an example, for the frequencies

¹See Jenkins and Watts (1968), p. 247.

$$f = \{0, .1, .2, .3, .4, .5\},$$

the corresponding periods are

$$\text{per.} = \{\infty, 10, 5, 3.3, 2.5, 2\}.$$

Therefore, to obtain the desired resolution for the longer frequencies, M will often have to be large; hence, the variance of the estimator will be large (a result of the large M).

The above discussion points to one very important reason why spectral analysis has not been used more extensively. Since one must often estimate a large number of frequencies to obtain the desired resolution, many data observations are needed. Otherwise, the variances of the spectral estimates will be so large as to preclude any confidence in them. For example, suppose data are available for 96 months and it is determined that $M = 48$ provides the desired resolution. Such an estimation, using the Parzen spectral window, produces only about seven degrees of freedom. The data problem becomes much worse when one attempts a multivariate spectral estimation.

The interpretation of the phase angle and associated leads and lags is a problem often ignored in studies using spectral analysis. The problem is that the phase angle is unique only in the range 0 to 2π (or equivalently, $-\pi$ to $+\pi$).¹ That is, any angle ϕ' , where

¹Actually, unless the analyst has prior information as to whether the two series are positively or negatively correlated, the phase angle is unique only in the range 0 to π (or $-\pi/2$ to $+\pi/2$).

$$\varphi' = \varphi \pm 2\pi k, \quad k \text{ any integer,}$$

will satisfy the definition of the phase angle. For example, a phase angle of 0.5π for a period of eight months could be interpreted as series x leading y by two months, or ten months, or even lagging y by six months. The lack of information about k causes the difficulty. Consequently, prior information must be included in the model.

A third problem that can cause difficulties in interpreting the spectral estimates in terms of an economic hypothesis is the lack of summary statistics. For example, suppose the relationship between x and y is estimated by cross-spectral techniques for 50 frequencies. Then the analyst must interpret the meaning of 50 points on the coherency spectrum, 50 points on the phase spectrum, and 50 points on the gain spectrum. He will be almost overwhelmed by estimates. Very often this will not cause large problems in cross-spectral analysis since he will be interested in only one part of the spectrum. However, this problem can become acute in multivariate spectral analysis since there are partial coherency spectra, a full coherence spectrum, partial phase spectra, and partial gain spectra. Therefore, although spectral analysis has the advantage of decomposing a time series into its frequency components, this decomposition has its drawbacks, for there are then no summary measures to describe the relationship.

From the discussion on the amount of data necessary for spectral estimation and the lack of summary statistics, it is clear that spectral analysis is limited in a practical sense to univariate and bivariate analyses most of the time. Occasionally, trivariate analysis may be used. Therefore, spectral analysis is precluded from many economics hypotheses just on the basis of the number of variables in the model.

IV. THREE EXAMPLES OF THE USEFULNESS OF SPECTRAL ANALYSIS

An important consideration in the application of spectral analysis to economics is whether or not the hypothesis can be reasonably stated in the frequency domain. One of the best examples of a hypothesis falling into this category is given by Scully (1971). In fact, Scully's hypothesis is probably much better stated in the frequency domain than in the time domain. Briefly, he attempted to examine the relationship between general business activity and strike activity. The hypothesis proposed by others was that strike activity should be directly related to, and a result of, general business activity. When business activity is high, there should be a larger number of work stoppages; when business activity is low, there should be few work stoppages.

This hypothesis has a very natural interpretation in the frequency domain, largely because general business activity is presumed to behave in a cyclical fashion. Scully's results point even further to the appropriateness of the frequency-domain statement of the model. He finds, as expected, that most of the variation in general business activity is explained by the lower frequency bands, consistent with the hypothesis that general business activity can be largely explained by long-run and cyclical components. On the other hand, he finds the somewhat surprising result that strike activity is largely explained by much shorter cycles -- on the order of 12 months. Therefore, the initial hypothesis that strike activity is a result of general business activity is refuted by this large discrepancy between the power spectra of the two time series. His results are further enhanced by the low coherence between the two series.

Therefore, spectral analysis represents the logical way to proceed to estimate Scully's hypothesis. However, sometimes one does not have such a natural frequency-domain hypothesis, even though it might best be examined in the frequency domain. That is, sometimes it is necessary to reformulate the hypothesis in the frequency domain although the hypothesis is logically stated in the time domain. An example of this is given by Cooper (1971 and 1972). Cooper's hypothesis represents a case where differing behavior is expected a priori across different frequencies.

Cooper examines the relationship between the money supply and stock returns. Through the combination of two different hypothesis -- the efficient capital markets model and the quantity theory of money -- it is expected that stock returns will lead the money supply in the longer run but will lag the money supply in the short run. That is, short-run deviations in money supply changes from the long-run trends and cycles will lead stock returns, and the long-run trends and cycles in money supply changes will lag stock returns. In addition, the combination of these two theories predicts that the relationship between money supply changes and stock returns is much stronger for the longer-run trends than for short-run deviations. These problems are further compounded by the fact that stock returns are expected to have a flat power spectrum -- that is, each frequency band contributes equally to the variance in the series. Therefore, since time-domain procedures average all frequency bands, one would expect to make a serious underestimate of the importance of the long-term relationship and possibly bias the coefficient estimates. This presents the extreme

case of errors in the variables, for the amount of variance explained by the primary frequencies is no larger than the variance explained by the other frequencies.

Time series regression analysis bears out the above predictions. Regressions for monthly data show returns leading money by roughly two months, thus indicating the dominance of the longer-run hypothesis. However, the R^2 for monthly data is on the order of 0.05. On the other hand, regressions for annual data show no lead or lag and an R^2 of more than 0.5. Therefore, the regression model does not permit adequate assessment of the hypothesis. The spectral estimation reveals quite a bit more. The coherence between money and returns was high for the lower one-sixth of the spectrum (on the order of 0.5 to 0.7) and quite low elsewhere, indicating the strength of the relationship for the low frequency bands. The phase shift indicates that money lags returns for the lower frequencies (by approximately one to three months) and leads for the higher frequencies (although the lead in the higher frequencies is subject to question because of the low coherence). Finally, although the regression coefficient for money is about 2.0, the gain measure for the lower frequencies is between 3.0 and 5.0, indicating the bias that may have occurred in the regression model.

The reason spectral analysis performs better than regression analysis in this case is a result of (1) the expected differing behavior across frequencies and (2) the flat spectrum for returns (the portions of the spectrum where there is expected to be little relation are averaged equally in the time domain representation).

The third hypothesis illustrates how errors in the variables affect the time-domain representation and how these might be interpreted in the frequency domain. Friedman's (1957) permanent income hypothesis indicates that income and consumption may each be decomposed into permanent and transitory components. Permanent consumption is expected to be a function of permanent income; transitory income and consumption are assumed to be unrelated. Therefore, a model relating measured income to measured consumption is subject to errors in the variables. The frequency-domain interpretation of this model would have the longer-run components of income and consumption related (corresponding to the "permanent" parts of the measured variables) and the shorter-run components unrelated. Since there are errors in the variables, one would expect that the marginal propensity to consume would be underestimated using ordinary regression techniques on the measured variables. Therefore, spectral analysis may present an attractive alternative for at least testing the validity of the hypothesis. If the coherence and gain are constant across frequencies, then one would reject the permanent income hypothesis (see Fishman [1970]). Alternatively, if the investigator can specify the model in the time domain (for example, suppose that "permanent income" is a geometrically declining weighted average of past values of measured income), then time domain techniques are useful. However, spectral analysis would be useful to test whether or not the permanent income hypothesis is valid without specifying the permanent component.

These three examples present cases where spectral analysis is a useful tool in economics. However, they are not intended to indicate that all hypotheses are appropriate for spectral analysis. Quite the contrary, the restrictive conditions developed in Section III indicate that spectral analysis is useful for only a few economic hypotheses.

V. CONCLUSIONS

Although spectral analysis offers an interesting way to estimate some economic hypotheses, it does not appear to be generally useful for economic problems as a whole. First, the hypothesis under investigation must have a reasonable interpretation in the frequency domain. Second, hypotheses that indicate substantial difference in behavior across frequencies may be amenable to spectral analysis. This is even more true if the frequency bands of interest explain only a small fraction of the total variance in the time series.

In addition to the theoretical restrictions, several practical elements limit the usefulness of spectral analysis. First, large data sets are generally required in order to make adequate estimates of the spectrum. Second, the lack of summary statistics often makes spectral analysis difficult to interpret. Third, spectral techniques are probably limited for the most part to univariate and bivariate models, with some possibility of extending to trivariate models. Therefore, many economic hypotheses are precluded from estimation in the frequency domain just because they cannot be represented as a simple two or three variable model.

One advantage to using spectral analysis is that the time series is broken down into individual frequency components, which can themselves be studied independently or together. In covariance analysis (and in regression analysis) the contributions of the different frequencies are all lumped together. However, this very advantage of spectral analysis is also its nemesis, for it is confusing to simultaneously examine and

evaluate the large number of frequency points.¹ For applications in economics, the estimates should be studied in groups (or frequency bands). The user should avoid both general summary statistics of the entire spectrum² (unless, of course, the estimates lend themselves to that type of measure) and frequency by frequency analysis of the spectrum (which may be of use in physical science applications). Also, in light of the large number of estimates, the user should attempt to evaluate the results in the framework of one or several different hypotheses, instead of merely presenting the various spectra. This is particularly true for the phase angle in cross or multivariate spectral analysis, for the phase angle has little meaning unless placed in the framework of some theory.

One aspect of the advantage of frequency domain analysis is the symmetric consideration given to both lags and leads between two series. For example, suppose that the low frequency components of $x(t)$ lead the low frequency components of $y(t)$ and that the high frequency components of $x(t)$ lag the high frequency components of $y(t)$. By its very nature, spectral analysis separates out these different effects, whereas time domain analysis may confuse these two distinct effects. That is, regression coefficients in time domain analysis lump together the effects of different frequencies. Therefore, it may be very difficult to sort

¹It is for this reason that inference is so difficult in spectral analysis. While the meaning of a confidence interval for a point estimate is clear, the meaning of a confidence interval for a function is not. Therefore, since the spectrum is a function of frequency, confidence bands cannot be used for the spectrum. One possible approach is to prescribe a loss function and to integrate the losses over all frequencies. This would yield a scalar loss rather than a vector of losses.

²Covariance and/or regression estimates provide such summaries.

out both leads and lags in the time domain representation. Alternatively, consider the following inventory problem. Suppose that firms effectively anticipate seasonal fluctuations in demand with their inventories but do not effectively anticipate cyclical or random movements in demand. Then, for seasonal frequencies inventories will lead sales while sales will lead inventories for nonseasonal frequencies. Spectral analysis will sort out this relation while the time domain results may not allow the analyst to distinguish these separate effects.

If used properly, spectral analysis can be a useful tool in economics. The real obstacle to the use of spectra lies not in the technical end (e.g., lag or spectral windows, truncation points, etc.) but in the interpretation of the spectral estimates.

REFERENCES

1. Cooper, Richard V.L., "Money and Stock Returns," Ph.D. Dissertation, University of Chicago, 1971.
2. _____, Efficient Capital Markets and the Quantity Theory of Money, The Rand Corporation, P-4886, August 1972.
3. Fishman, George, Spectral Methods in Econometrics, Harvard University Press, Cambridge, Massachusetts, 1969.
4. Friedman, Milton, A Theory of the Consumption Function, Princeton University Press, Princeton, New Jersey, 1967.
5. Granger, C.W.J., "The Typical Spectral Shape of an Economic Variable," Econometrica, 34, January 1966, pp. 150-161.
6. Hannan, Edward J., "Regression for Time Series," in M. Rosenblatt (ed.), Time Series Analysis, John Wiley and Sons, New York, 1963.
7. _____, "The Estimation of Relationships Involving Distributed Lags," Econometrica, 33, January 1965, pp. 206-224.
8. _____, Multiple Time Series, John Wiley and Sons, New York, 1970.
9. Hause, John C., "Spectral Analysis and the Detection of Lead-Lag Relationships," American Economic Review, 61, March 1971, pp. 213-217.
10. Jenkins, Gwilym M., and Donald G. Watts, Spectral Analysis and its Applications, Holden-Day, Inc., San Francisco, 1968.
11. Nerlove, Marc., "Spectral Analysis of Seasonal Adjustment Procedures," Econometrica, 32, July 1964, pp. 241-286.
12. Scully, Gerald W., "Business Cycles and Industrial Strike Activity," Journal of Business, 44, October 1971, pp. 359-374.